arXiv:1401.3816v1 [physics.flu-dyn] 16 Jan 2014

Physics-Based Learning Models for Ship Hydrodynamics

G.D. Weymouth[a], Dick K.P. Yue[a]

[a] *Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, 02139, USA*

Abstract

We present the concepts of physics-based learning models (PBLM) and their relevance and application to the field of ship hydrodynamics. The utility of physics-based learning is motivated by contrasting generic learning models (GLM) for regression predictions which do not presume any knowledge of the system other than the training data provided, with methods such as semi-empirical models which incorporate physical-insights along with data-fitting. PBLM provides a framework wherein intermediate models (IM), which capture (some) physical aspects of the problem, are incorporated into modern GLM tools to substantially improve the predictions of the latter, minimizing the reliance on costly experimental measurements or high-resolution high-fidelity numerical solutions. To illustrate the versatility and efficacy of PBLM, we present three wave-ship interaction problems: (i) at speed waterline profiles, (ii) ship motions in head seas, and (iii) three-dimensional breaking bow waves. PBLM is shown to be robust and produce error rates at or below the uncertainty in the generated data, at a small fraction of the expense of high-resolution numerical predictions.

Keywords: Computers in Design, Hull Form Hydrodynamics, Seakeeping, Machine Learning, CFD

1. Introduction

The state of the art in computational fluid dynamics (CFD) has advanced rapidly in the last decades, complementing advances in whole field quantitative experimental measurements. Advanced numerical methods and high performance computing capabilities have produced high fidelity simulations of ship hydrodynamics such as time domain predictions of the motions and flow around a fully appended ship, resolving details of the breaking bow and stern waves and the underlying turbulent flow (e.g. Weymouth et al., 2007; Carrica et al., 2007). Such first-principle physics-based tools achieve accuracy through minimization of numerical and user errors. Despite increases in computing resources, these methods are still extremely computationally expensive (requiring hundreds of millions of grid points and hundreds of thousands of CPU hours), and generally still depend on highly educated and experienced users. As such, they are still mainly appropriate only for final stages of the design process.

Email addresses: weymouth@mit.edu (G.D. Weymouth), yue@mit.edu (Dick K.P. Yue)

Learning-based models offer an attractive alternative to first-principle physics-based simulations. Such models can be extremely fast, with design evaluations often requiring less than a second on personal computers. In contrast to first-principle methods, learning models require an available set of example (or training) data, which we denote by: $\{y_i|\mathbf{x}_i\}$ for $i = 1 \dots n$, where y is the dependant variable to be estimated (e.g., drag, heave motion, etc) and $\mathbf{x}$ is the vector of independent variables (e.g., ship speed, incident wave frequency, ...). This data is typically obtained from experiments or high-fidelity simulations. The learning model seeks to approximate the unknown function f from which the example data are generated: $y = f(\mathbf{x}) + \epsilon$, where ϵ is the random error in the measured/example data. The learning model $f_{\mathcal{G}}$ approximating f:

$$f_{\mathcal{G}}(\mathbf{x}|\mathbf{c}) \approx f(\mathbf{x}) \tag{1}$$

is 'trained' by optimizing internal coefficients, $\mathbf{c}$, to minimize error on the training data. The predictive error of the model is then tested on points held out of the training set.

In this work, we will use the name generic learning model (GLM) for models $f_{\mathcal{G}}$ which approximate the unknown function f using no system-specific information other than the training set. There is a vast literature for GLMs, from basic least-squares polynomials fits to splines to non-linear multi-layer neural networks. A common limitation of all GLMs is that they can only make accurate test predictions when available training data completely maps out the test space. Regions of the design space with few example data have a high risk of prediction errors, and filling in the gaps requires costly experiments or computations. In addition, more complex physical systems typically necessitate more elaborate GLMs which must be trained with correspondingly larger sets of examples (Evgeniou et al., 2000). This can also lead to over-fitting; defined as when an overly flexible GLM is used to reduce the training error, only to find poor generalization to the unseen test data (Vapnik, 1995; Burges, 1998; Girosi et al., 1995). Utilizing techniques such as cross validation (Golub et al., 1979) and regularization can help minimize over-fitting (and will be exploited herein) but they do not address the systemic concern: a generic learning model does not know anything about the system other than the data it is provided.

This deficiency in GLMs is a well established historical problem and significant effort has been dedicated to the task of relieving the data dependence of such regression models through the addition of some level of 'physical insight'. As an example, consider the scaling of the resistance of a ship proposed in Froude (1955). In modern terminology, Froude's insight was that a ship's resistance is a function of the Reynolds number Re and the Froude number Fr and that these factors act independently. In other words, the total resistance coefficient C_T can be approximated by an additive model

$$C_T = f(Re, Fr, \mathbf{x}_{geo}) \approx C_r(Fr, \mathbf{x}_{geo}) + C_f(R_e)$$

where C_f, the friction coefficient is assumed to be a function only of Re, and C_r, the residuary resistance coefficient is a function only of Fr (and the ship geometry represented by input geometry parameters $\mathbf{x}_{geo}$). The decomposition of the Re dependence is a physical insight which enables the friction model to be determined independently using flat plate data, thereby allowing data from Fr-scale model tests to be used in predictions of full-scale ship resistance.

The 1957 ITTC semi-empirical model for the friction coefficient C_f is another example of a data-based model incorporating physical insights obtained from turbulent flat-plate boundary layer (Granville, 1977):

$$C_f(Re|\mathbf{c}) = \frac{c_1}{(\log Re - c_2)^2}. \tag{2}$$

In (2), $\mathbf{c} = \{c_1, c_2\}^T$ are the internal model coefficients, and setting $c_1 \approx 2$ and $c_2 \approx 0.075$ result in model predictions which agree with the experimental resistance data taken from a fully submerged flat plate to within 9% over a wide range in the input ($Re = 10^5 \sim 10^{10}$). In comparison, a generic model function $f_{\mathcal{G}}$ in polynomial form with the same number of free parameters

$$f_{\mathcal{G}}(Re|\mathbf{c}) = c_1 + c_2 Re$$

and a least-squares fit leaves regions of the data with more than 20% error. The simple semi-empirical model (2) illustrates how physics-based insights in a learning (regression) model can greatly increase predictive accuracy without increasing the number of free parameters.

Extending this general approach to predict the behavior of systems which are too complex to be modeled by an analytic semi-empirical model is an active research topic. Successful examples in the naval architecture literature are the ONR ship and submarine maneuvering prediction models of Hess et al. (2006); Faller et al. (2005, 2006). Unlike time-domain unsteady CFD ship motions predictions (such as Weymouth et al. (2005); Carrica et al. (2007)), these models feed pre-computed steady-state force predictions (along with other system variables such as relative speed and bearing) into a recursive neural network learning model to generate predictions of the resulting ship motions in real-time. Such research illustrates that a combination of a GLM (the recursive neural network) and a simplified physics-based model (the steady-state force predictions) with high-fidelity data can produce useful design tools. The ONR learning model has a complex structure which has been specially designed for maneuvering predictions. The question we seek to answer is whether there is a general framework for such ideas that can be applied to a variety of ship hydrodynamics and ocean engineering problems, taking advantage of physics-based insights that might be obtained from preliminary design tools such as potential flow models.

In this work we develop a general approach for physics-based learning models (PBLM) which extends historical approaches such as analytic semi-empirical models and expands upon modern developments such as the ONR maneuvering predictions. The key idea behind PBLM is the robust incorporation of one or more physics-based intermediate models (IM) within a GLM. Rather than increasing the complexity of $f_{\mathcal{G}}$ to improve the fit, PBLM complements a simple GLM with expanded basis functions constructed from fast and robust physics-based IM(s). By complementing the GLM and the training data with physical behaviors (and constraints) inherent in the IM, we demonstrate that PBLM obtains significantly improved accuracy while reducing data dependence and the risk of over-fitting.

We develop our PBLM function $f_{\mathcal{P}}$ from three primary components: a training data set $\{y|\mathbf{x}\}$; a generic learning model (GLM) defined by $f_{\mathcal{G}}$; and a physics-based intermediate model (IM) defined by $f_{\mathcal{I}}$. (While the methodology easily generalizes to the use of more than one IM, for clarity and simplicity we assume hereafter only one IM.).

The PBLM framework applies to many simple and robust GLMs, but to illustrate the power of PBLM, we use regularized non-parametric GLMs, which are explained in §2.1, after covering some fundamental topics in machine learning. In §2.2, our method for constructing $f_{\mathcal{P}}$ from these components is developed using a simple example of the transient development of a body in resonant motion. To demonstrate the usefulness and performance of PBLM to realistic ship hydrodynamics applications, we apply this new framework to three wave-ship interaction problems: (i) at speed waterline elevation profiles over a range of Froude numbers (§3.1), (ii) pitch and heave response of a ship with forward speed (§3.2), and (iii) breaking bow wave characteristic predictions for variable ship speed and breadth (§3.3). The first two cases utilize the Wigley hull and are chosen because of their multi-dimensional input space and the wealth of classical experimental data for this canonical hull form. The last case introduces geometry dependence, nonlinear wave mechanics, and uses computational sources for both the training data and the physics-based intermediate model. In all cases, we show that the PBLM approach achieves accuracy levels equivalent to high-fidelity experimental or CFD predictions, are significantly less data dependant than generic learning models, and have very small overall prediction costs.

There are many other branches of numerics which are related to the current work including the field of data assimilation using ensemble Kalman filtering (Kalman, 1960; Evensen, 2003) and the field of multi-level optimization methods such as Approximation and Model Management Optimization (AMMO: Alexandrov et al. (1998, 2001)). The ensemble Kalman Filter (enKF) is widely used in many fields such as climate and weather forecasting which rely on simplified models. AMMO is used to in many types of engineering optimization problems where evaluating approximate objective functions leads to great computational saving.

The current work is influenced by these methods and others which use multiple levels of information to develop predictions, but the problem we have set out for ourselves is distinct. The the enKF is used to recursively filter data using large ensembles of simple forecasting models with randomized parameters. On the other hand, AMMO is utilized to establish trust regions for the approximate objective function to ensure convergence to the optimum of the true objective function. In contrast, the goal of the PBLM is to use an intermediate model with no free parameters to help generalize a small set of high fidelity examples to achieve accurate predictions across the input space. This could be seen as the sparse-data sparse-model limit of the filtering problem, or as an optimization problem which requires accurate predictions of the full response space. With these distinctions in mind, our development in the following section focuses on the literature of machine learning from labelled data as it is most directly applicable to our problem setting.

2. Methodology

To support the development of our PBLM approach, we first present some key concepts in machine learning. Consider a physical problem wherein we are given n data points $\{y_i|\mathbf{x}_i\}$ for $i = 1 \dots n$ generated by the unknown function f. The simplest generic model of this function would be a weighted sum of the inputs $\mathbf{x}$, ie

$$f_{\mathcal{G}}(\mathbf{x}|\mathbf{c}) = \mathbf{x}^T\mathbf{c} \tag{3}$$

which is linear in the inputs and the parameters $\mathbf{c}$. The model parameters are optimized to minimize an error metric $H(f_{\mathcal{G}})$ such as the sum of squared error over the n points in the training set:

$$H(f_{\mathcal{G}}) = ||y_i - f_{\mathcal{G}}(\mathbf{x}_i)||^2 = \sum_{i=1}^{n} \left[y_i - f_{\mathcal{G}}(\mathbf{x}_i)\right]^2 .$$

Unfortunately, a linear superposition of the input parameters (3) rarely provides an adequate fit to the data describing complex engineering systems. This can be improved by expanding the degrees of freedom using a parametric model

$$f_{\mathcal{G}}(\mathbf{x}|\mathbf{c}) = \boldsymbol{\gamma}(\mathbf{x})^T \mathbf{c} \tag{4}$$

where $\boldsymbol{\gamma}(\mathbf{x})$ are a vector of generic basis functions, such as a polynomial basis $(1, x, x^2, \ldots)^T$. While the form of (4) is simple, it still leaves the complex matter of defining the basis set $\boldsymbol{\gamma}$. Using a low-dimensional basis with too few degrees of freedom will not help the fit, and least-squares fitting of high-dimensional basis are prone to over-fitting and can become numerically unstable.

Our main contribution to methodology is to introduce a physically motivated basis via a physics-based intermediate model (IM) to increase accuracy without over-fitting. This PBLM formulation is developed in §2.2, but first we consider a generic learning model solution, non-parametric regularized learning.

2.1. Generic learning models using non-parametric regularization

We introduce two generic machine learning approaches to help reduce training error without over-fitting; non-parametric regression and regularization. We discuss these methods because they introduce concepts such as effective degrees of freedom, and form the basis of the GLMs used in our PBLM approach.

Non-parametric regression methods such as splines and kernel interpolation ensure a good fit to the training data by generating a generic basis function for each data point. Typically non-parametric methods use a generic learning function of the form

$$f_{\mathcal{G}}(\mathbf{x}|\mathbf{c}) = \sum_{i=1}^{n} c_i G(\mathbf{x} - \mathbf{x}_i) = \sum_{i=1}^{n} c_i \gamma_i(\mathbf{x}) \tag{5}$$

where G is a symmetric kernel function which is used to generate a generic basis functions γ_i for every data point $\mathbf{x}_i$. Unlike a polynomial basis, these kernels are localized, e.g., a Gaussian kernel $G(\mathbf{x}) = \exp(||\mathbf{x}||^2/\sigma)$, limiting the influence of each basis function to a subregion of the input space. This enables the non-parametric model to accurately fit the training data from a wide variety of systems. However, non-parametric models have up to n degrees of freedom which makes them very prone to over-fitting.

Regularization (also known as ridge regression) reduces over-fitting by introducing a 'statistical-insight' into the GLM, namely: smooth predictions tend to generalize better to unseen data. To regularize a GLM, the parameters are optimized by minimizing a regularized error metric:

$$H(f_{\mathcal{G}}) = (1 - \lambda)||y_i - f_{\mathcal{G}}(\mathbf{x}_i)||^2 + \lambda J(f_{\mathcal{G}}), \tag{6}$$

where the regularization function $J(f_{\mathcal{G}})$ is a term which penalizes the higher derivatives of $f_{\mathcal{G}}$ and λ is the regularization constant bounded by $0 \leq \lambda < 1$. The effect of regularized learning is a reduction in the effective degrees of freedom, eDOF, of the GLM, given by

$$\text{eDOF} = \sum_{i=1}^{n} \frac{(1-\lambda)d_i^2}{(1-\lambda)d_i^2 + \lambda} \tag{7}$$

where d_i are the eigenvalues of the linear system used to minimize Eq. 6.

The parameter λ determines the trade-off between data-fit and model smoothness. Setting $\lambda = 0$ results in an unconstrained least squares model with eDOF=n. This model would be flexible, but would over-fit the n training data points. When $\lambda \to 1$, the function $f_{\mathcal{G}}$ is constrained to be linear, with eDOF$\to 0$. This model would be robust, but is not likely to predict the system behaviour accurately. An intermediate value of λ is determined automatically as in Golub et al. (1979) by minimizing the generalized cross validation error on the training set in an attempt to balance the desire for high eDOF to fit the training data and the desire for low eDOF to smooth $f_{\mathcal{G}}$ and improve generalization.

We use three well-established versions of regularized non-parametric regression in our applications below: thin plate regression splines (Wood, 2003, 2008), gaussian regularization networks (Girosi et al., 1995; Burges, 1998), and support vector machines (Vapnik, 1995; Evgeniou et al., 2000; Scholkopf et al., 2000). In all cases, the regularized problem is more numerically stable than least-squares and has a unique solution for the optimum parameters, unlike non-linear methods such as neural networks. However, the model performance is still dependant on the amount of training data and is highly sensitive to the chosen value of λ (and any parameters in the kernel function G).

2.2. Physics-Based Learning Models

The regularized non-parametric GLMs presented in the preceding section are flexible and offer a level of protection from over-fitting but they are still generic methods and cannot make reliable predictions in data-sparse regions of the input space. We will reduce this data dependence not by fine-tuning parameters or adopting more complex GLMs but by enhancing these robust learning tools with physics-based information from an intermediate model $f_{\mathcal{I}}$. Additionally, instead of trying to develop a new learning approach for each new system, we develop a general PBLM approach wherein the physics-based learning model $f_{\mathcal{P}}$ is of the form

$$f_{\mathcal{P}}(\mathbf{x}) = f_{\mathcal{G}}(\mathbf{x}, \boldsymbol{\rho}(\mathbf{x})), \tag{8}$$

augmenting a simple GLM with a physics-based basis $\boldsymbol{\rho}$ to reduce data-dependence and increase accuracy.

To illustrate the basic ideas and efficacy of PBLM, consider a simple model problem of the transient startup from rest of the resonant harmonic motions, say corresponding to initial development of the resonant response of a floating body in regular waves. For simplicity, we assume the true motion is given by

$$f(t) = \alpha \tanh(\beta t) \sin(\omega t + \phi) \tag{9}$$

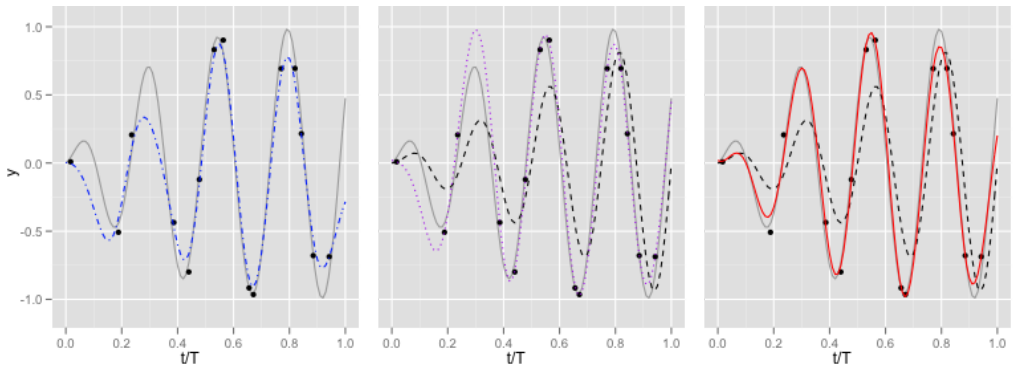

(a) Data-based GLM (b) IM and additive model (c) PBLM

Figure 1: PBLM example for the transient startup of harmonic motion. In all figures — is the true system function f and • are the n=15 randomly sampled values of f with Gaussian noise of σ=10%. In figure (a), – · – is the optimized Gaussian regularization network GLM fit. In figure (b), the true function f is compared to −− the simple IM $f_{\mathcal{I}}$ (11), and · · · the predictions of the additive model $f_{\mathcal{A}}$ (12). In figure (c) — is the prediction of PBLM (13) using the phase-shifted basis $\boldsymbol{\rho}$. The PBLM results are more than twice as accurate than either the GLM or the additive model, despite the small training sample and the fact that the IM has 32% error.

where β is the startup time scale, ω is the frequency of the regular incident wave, and ϕ phase difference between the response and the wave; and that we have a limited sample of experimental data of this system

$$y_i = f(t_i) + \mathcal{N}(\sigma), \quad i = 1 \dots n, \tag{10}$$

at n randomly sampled times, t_i, where the $\mathcal{N}$ term is a zero-mean Gaussian noise in the data (see Fig. 1).

The first approach is to fit (10) with a GLM. In this case, linear and parametric GLMs give universally inadequate predictions because of the high frequency and variance in the system response. Non-parametric GLM results are highly dependant on the form of the kernel G, the regularization parameter λ, and the exact data sample provided. Fig. 1(a) shows the predictions of an optimized Gaussian regularization network GLM. The result are representative of a well-tuned model, with the fit being better in regions with more data and worse in region with less, with an overall RMS error of 20%. The predictions, however, are not robust, with high sensitivity to changes in the training data and parameter values.

To improve the performance of the GLM for this sparsely sampled system, we introduce a simple IM which captures merely the physics-based expectation that (a) in steady state, the response should be harmonic with frequency ω; and (b) the initial resonant response grows linearly in time. Without additional knowledge of the motion amplitude or phase, or the time required to reach steady state, we construct an IM simply as, say,

$$f_{\mathcal{I}}(t) = t \sin(\omega t). \tag{11}$$

This function, after appropriate scaling, is plotted in Fig. 1(b) and while qualitative similarity to the true response, it has a quantitative root-mean square error (RMSE) of over 30%.

While (11) has low predictive accuracy, it encapsulates additional knowledge of our system and can be used to improve the GLM prediction by supplementing the data with this information. In this simple example, we could parametrize $f_{\mathcal{I}}$, turning it into a semi-empirical model. However, this approach does not apply to general intermediate models. The simplest way to use the IM as a 'black-box' is to treat the difference between $f_{\mathcal{I}}$ and the data as a nuisance variable and fit that error with a GLM. Taking this idea one step further, we can use a linear additive model (Schölkopf and Smola, 2001; Wood, 2008)

$$f_{\mathcal{A}}(t|\mathbf{c}) = f_{\mathcal{I}}(t)c_\rho + f_{\mathcal{G}}(t|\mathbf{c}_\gamma) \tag{12}$$

which linearly weights the IM predictions with GLM predictions to obtain the best overall fit to the data. The additive model is an improvement over the nuisance model because the physics-based coefficient c_ρ is determined simultaneously with the generic coefficients $\mathbf{c}_\gamma$. The nuisance and additive models incorporate information from intermediate models in a general way, but they will only improve upon a GLM when the IM is a very good predictor of the system. Fig. 1(b) shows that the additive model predictions for this example system are only marginally better than the GLM results, with the RMSE near 16%. Additionally, the performance is still highly dependant on the details of the GLM and the exact data sampled.

While the additive model is a step in the right direction, it is too reliant on the GLM and the accuracy of the function $f_{\mathcal{I}}$. Our solution is to expand the physics-based degrees of freedom into a representative basis using a Taylor expansion of $f_{\mathcal{I}}$ (to provide gradient information). This gives the learning model more information about the form of $f_{\mathcal{I}}$ and more flexibility in constructing a physics-based solution before relying on the generic basis $\boldsymbol{\gamma}$. Development and evaluation of an exact Taylor expansion for general IMs would be a significant expense. However, one of the central advantages of shifting our information burden from the data to the IM is that the IM can be evaluated anywhere very quickly, not just at the x points where the data is available. Therefore, we can establish properties of the IM, such as the local derivatives, using a finite difference approximation

$$\boldsymbol{\rho}(x) = \left(f_{\mathcal{I}}(x), \frac{f_{\mathcal{I}}(x+\Delta x) - f_{\mathcal{I}}(x-\Delta x)}{2\Delta x}, \frac{f_{\mathcal{I}}(x+\Delta x) - 2f_{\mathcal{I}}(x) + f_{\mathcal{I}}(x-\Delta x)}{\Delta x^2} \dots \right)^T.$$

where the points $x \pm \Delta x$ are not required to be in the set of points x for which we have training data. Through linear combinations, the finite difference is equivalent to

$$\boldsymbol{\rho}(\mathbf{x}) = (f_{\mathcal{I}}(\mathbf{x}), f_{\mathcal{I}}(\mathbf{x}+\Delta\mathbf{x}), f_{\mathcal{I}}(\mathbf{x}-\Delta\mathbf{x}))^T. \tag{13}$$

This phase-shifted physical basis is equivalent to the derivative basis in the limit of $\Delta x \to 0$, is easily computable using the intermediate model $f_{\mathcal{I}}$, and enables the generation of robust and accurate PBLMs. In cases where phase-shifting either left or right of a data point moves past a physical boundary (such as zero speed) the offending term is simply left out and the implied derivative reverts to a one-sided first-order approximation.

We develop two simple and robust methods of including the phase-shifted physical basis $\boldsymbol{\rho}$ in our PBLM. First, we can construct a parametric additive model

$$f_{\mathcal{P}}(\mathbf{x}|\mathbf{c}) = \boldsymbol{\rho}(\mathbf{x})^T\mathbf{c}_\rho + f_{\mathcal{G}}(\mathbf{x}|\mathbf{c}_\gamma) \tag{14}$$

which complements the generic basis $\boldsymbol{\gamma}$ with the physical basis $\boldsymbol{\rho}$. A non-parametric version of the form of (5) can also be constructed using $\boldsymbol{\rho}$ to complement the input space, i.e.,

$$f_{\mathcal{P}}(\mathbf{x}|\mathbf{c}) = f_{\mathcal{G}}(\mathbf{z}|\mathbf{c}) = \sum_{i=1}^{n} c_i G(\boldsymbol{\rho}(\mathbf{x}) - \boldsymbol{\rho}(\mathbf{x}_i)) G(\mathbf{x} - \mathbf{x}_i) \tag{15}$$

where $\mathbf{z} \equiv (\mathbf{x}^T, \boldsymbol{\rho}(\mathbf{x})^T)^T$. This weighted kernel formulation introduces physical-relevance into the PBLM while utilizing the advantages of localized non-parametric regression. The predictions using Eq. 15 are shown in Fig. 1(c). The accuracy is excellent, with an RMSE of 8%, lower than the noise in the data and less than half of the GLM error. Additionally, this performance level is robust to changes in the details of the GLM and the training sample.

This simple one dimensional example turns out to be quite representative of the later results for realistic complex problems. Furthermore, our PBLM construction above is completely general, requiring no adjustment or modifications to produce predictions for real ocean engineering systems with multidimensional input spaces, as illustrated in the later applications in §3.

We have developed both (14) and (15) to demonstrate that physics-based learning is a general concept which may be used in different specific machine learning models. The parametric form of (14) is simpler, and more familiar to anyone who has used a linear solver. The non-parametric form of (15) is more flexible, but it requires a bit of additional machinery to implement. Additionally, different types of learning problems may dictate different model structures. Parametric models can be more readily interpreted by examining how the basis functions have been weighted, whereas problems with Runge's phenomenon may be avoided using a non-parametric approach. We will use (14) in our first example in §3.1 and (15) is used in our examples in §3.2 and 3.3.

Before moving to the applications, we remark on the desired attributes of the IM used to construct the physical basis $\boldsymbol{\rho}$. As in the example above, the ideal $f_{\mathcal{I}}$ should be efficient, stable, robust, provide good descriptions or approximations of (some aspects of) the system, and does not itself depend on free parameters or require new system-specific learning or user artistry. Note that the inclusion of an intermediate model function $f_{\mathcal{I}}$ which is a poor approximation of the system, while not ideal, is also not catastrophic. An inaccurate IM will be poorly correlated to the training data, and will be ignored by the learning machine (given no weighting) meaning the PBLM will simply preform as a GLM. An illustration of this is shown in the example in §3.2. A more serious constraint is that the IM prediction be noise free. As the IM is used to construct a basis, any noise in $f_{\mathcal{I}}$ would propagate up to $f_{\mathcal{P}}$ giving undesirable noisy final predictions. Practically, this can almost always be achieved by filtering away background noise in the IMs before using them in a PBLM as demonstrated in the last example problem of §3.3.

3. Applications

To illustrate the application and performance of the PBLM method developed above, we consider in this section three ship hydrodynamics applications involving the prediction of (i) the waterline profiles on a ship over a range of Froude numbers, (ii) the heave and pitch responses of a ship advancing in regular seas, and (iii) the breaking bow waves for a ship of varying forward speed and beam-to-length ratio.

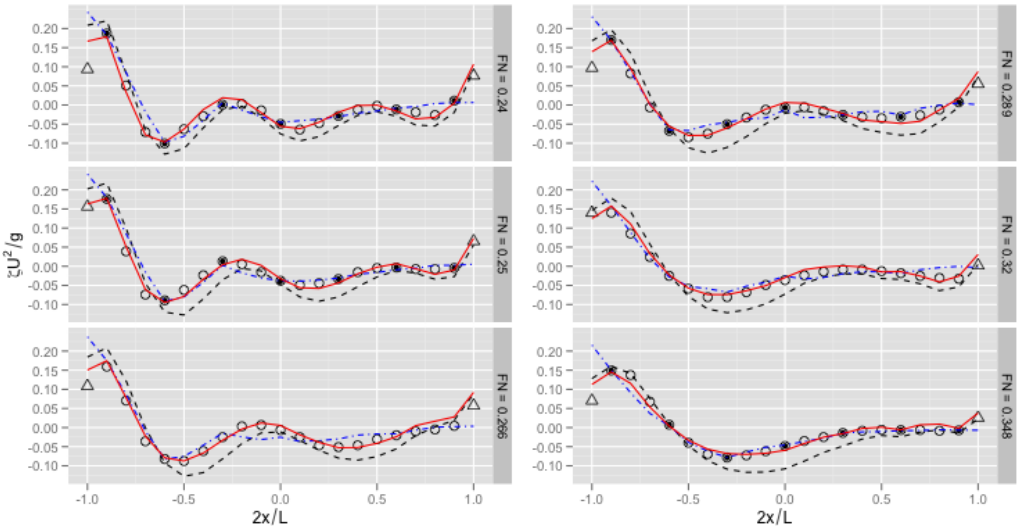

Figure 2: Waterline elevation profile data and predictions for the Wigley hull over six Froude numbers. Symbols denote the experimental measurements: • are the training data, ○ are the test data, and △ are extrapolation points outside the training data. Lines denote the prediction methods: – – is the potential flow IM prediction, – · – is the GLM prediction (using thin plate regression spline), and — the PBLM prediction. PBLM obtains significantly increased accuracy away from training data, interpolating accurately to the unseen Froude numbers, and extrapolating to recover physically reasonable predictions for the bow and stern elevations.

3.1. Waterline Predictions for the Wigley Hull

We consider the predictions of the waterline on a Wigley hull over a broad range of Froude numbers. These classical experiments have served as benchmark data for many numerical methods including advanced potential flow and Navier-Stokes simulations. Even for this case of steady flow and relatively smooth profiles, a generic learning model trained using experimental data set has large errors in data-sparse regions and makes non-physical predictions. We demonstrate that by complementing the GLM with a simple linear potential flow intermediate model, PBLM achieves excellent predictive accuracy across the input space, even at untrained Froude numbers.

The data used for training and testing our models comes from the experimental measurements of Shearer and Cross (1965) which reports the at-speed waterline profile for a fixed sinkage and trim Wigley hull. From this source we select data for the normalized waterline elevation $\eta U^2/g$ over six Fr speeds and 21 $2x/L$ section positions. These waterline profile measurement points are plotted in Fig. 2, showing the dependence of the profiles on Froude number.

For the GLM, we use thin plate regression splines of Wood (2003), a regularized multi-dimensional spline method. We train this generic learning model to predict the waterline profile over the two-dimensional input vector by minimizing the regularized error (Eq. 6) over the training set $\{y = \eta U^2/g \mid \mathbf{x} = (Fr, 2x/L)^T\}$. A third of the profile points were included in the training data except for two intermediate Froude numbers (Fr=0.266, 0.32 in Fig. 2) which were completely held out to test the model's ability to generalize to unseen speeds. The held out points within the convex hull of the

	IM (potential flow)	GLM	PBLM
Test set RMSE	0.193	0.201	0.085
Total RMSE	0.187	0.177	0.078
eDOF	0	21.2	5.0

Table 1: RMSE and effective degrees of freedom (eDOF, Eq.(7)) for the prediction of the Wigley hull waterline elevation. The RMSE, scaled by the maximum response $\eta U^2/g = 0.187$, are given for the test set as well as for the total (including the error on the training, test, and extrapolation points) compared to tank measurements. The GLM is the thin plate regression spline, and the IM is simple potential flow with reflection free-surface boundary condition. The RMSE of the IM and GLM are comparable (close to 20%), while the PBLM incorporating these (using the phased-shifted physics-based basis with only 5 effective degrees of freedom) is able to reduce the test and total errors but a significant factor.

training data make up the test set, used to determine the generalized accuracy of the GLM predictions. Predictions of the elevation near the bow and stern are extrapolations outside the training data, making these especially difficult for the GLM. With the model trained, the GLM predicts the wave elevation at every point in the parameter space using trivial computational effort. Considering the sparsity of the supplied training data set, the thin plate regression spline GLM performs fairly well as summarized in Tab. 1. The test RMSE is around 20% of the maximum waterline elevation measurement. However, as expected of any generic learning model, the error increases for predictions far from the training data, for example in the unseen Froude numbers. The extrapolated predictions for the leading and trailing edge wave heights are also non-physical, missing the low initial wave height on the bow and failing to predict the run up at the stern.

For the intermediate model we use a first-order potential flow panel method and a reflection (double-hull) boundary condition for the free surface as in the classical approach of Hess and Smith (1964). This simple model is chosen for the speed and robustness of the IM predictions, rather than their accuracy. All of the potential flow numerical parameters (related to the panelization, inversion etc.) are fixed to avoid propagation of learning capacity to the PBLM. The waterline predictions using this IM are plotted in in Fig. 2. As summarized in Tab. 1, these predictions have an RMSE similar to the GLM, around 20%.

For the PBLM, we construct the phase shifted basis $\boldsymbol{\rho}$ of Eq. 13 using the potential flow IM. The basis is incorporated into the GLM through the simple additive model of Eq. 14 and retrained on the training data. The predictions for the PBLM are shown as the red solid line in Fig. 2 and summarized in Tab. 1. The addition of $\boldsymbol{\rho}$ has reduced the error on the test data by 58% and the figure shows the accuracy is much less dependent on the proximity of the training data. The figure also shows that the predictions of the trailing edge height are excellent despite these points being extrapolations beyond the training data. The leading edge profile heights are generally over-predicted, but always physically realistic.

The effective degrees of freedom (eDOF) of GLM versus PBLM are given in Tab. 1 quantifying the decreased data dependence of the PBLM. The potential flow model has no adjustable degrees of freedom (eDOF=0) and the GLM has effectively 21 degrees of freedom after regularization. As there are only 28 points in the training set, the GLM is only weakly regularized, which translates to higher data dependence. In contrast, PBLM has very little data dependence, using only 5 degrees of freedom to fit the data more than

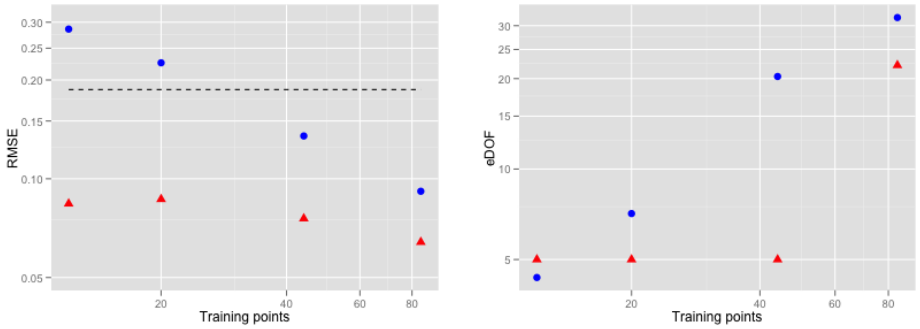

(a) (b)

Figure 3: Convergence study on the dependence of the test RMSE (a) and eDOF (b) on the number of training points for the GLM (●) and PBLM (▲). The RMSE is scaled by the maximum response $\eta U^2/g = 0.187$ and the potential flow RMSE is shown for reference (−−). The training sets were selected using systematic reduction, and in all cases the Fr=0.266,0.32 data have been held out. Note that the PBLM is essentially converged using 12 training data points (3 points per speed), and at this level outperforms the GLM using 84 points.

twice as accurately.

A systematic study further quantifies the dependence on the number of training points and is presented in Fig. 3. For this study, the number of waterline points used in the training set was systematically reduced using a 'point-skipping' approach; starting with all 21 point in each profile, then using every other point, and eventually ending up with only the points at $2x/L = -1, 0, 1$ in the training data. The points in each coarse set are contained in all the finer sets, as in a nested quadrature. In all cases the Fr=0.266,0.32 data have been held out of the training set. The results in Fig. 3 show that the GLM response is completely dependent on the number of training points, making steady gains in accuracy with increased data (and increased eDOF). In contrast, with as few as 12 training points the PBLM solution has essentially converged on the solution shown in Fig. 2 with RMSE≈8.5%, which is more accurate than the GLM solution using all 84 available waterline points. Only with the full set of points does the PBLM relax it's regularization and increase from 5 to 22 degrees of freedom to more closely fit the data, achieving 6.4% test RMSE. These results quantify that the addition of the physics-based basis $\boldsymbol{\rho}$ is remarkably effective at shifting the learning dependence from the expensive data measurements to the inexpensive intermediate model.

3.2. Pitch and Heave Motions Predictions

We next consider the prediction of the heave and pitch response amplitude operators (RAO) for the Wigley hull in head seas. As in the previous example, this problem has been used for the validation of many numerical prediction methods including the unsteady-RANS CFD predictions of Weymouth et al. (2007). The RANS predictions were found to be quite accurate with RMSE levels of around 2.5%, but each $(Fr, \lambda/L)$ evaluation took $O(10^4)$ CPU hours. We show that by incorporating a simple linear potential flow IM into a standard GLM we achieve a similar level of accuracy over the

Heave	IM (LAMP-I)	GLM	PBLM
Test set RMSE	0.047	0.042	0.025
Total RMSE	0.054	0.027	0.016
eDOF	0	16.4	10.3
Pitch	IM(LAMP-I)	GLM	PBLM
Test set RMSE	0.131	0.079	0.029
Total RMSE	0.141	0.053	0.030
eDOF	0	13.0	7.8

Table 2: Prediction errors for the test set and the total RMSE (which includes the error on the test and training points) and the effective degrees of freedom (eDOF) for the Wigley hull pitch and heave tests. The root mean square errors are scaled by the maximum response for each test ($Y_3/a = 3.54$, $Y_5 L/2\pi a = 2.0$). The PBLM demonstrates error levels below 3% (equivalent to the error level of time domain Navier-Stokes simulations) using simple linear models and extremely sparse training data.

complete input space as Weymouth et al. (2007) using a fraction of a second of computing time. We also find that the PBLM predictions are robust to changes in the training data.

The data for this problem is taken from the experimental measurements of Journee (1992) for the Wigley hull I and wave amplitude $a/L \approx 0.005$. For this example we use data for the normalized heave (Y_3/a) and pitch ($Y_5 L/2\pi a$) response over three Fr speeds and $\sim$11 λ/L head sea wavelengths. The data is denoted by the point symbols in Fig. 4 showing substantial variations in the RAO peak frequency and height with changing Froude number, reflecting the complex wave-ship-flow interactions at play.

For this example, we choose the Gaussian regularization network GLM which is studied extensively in machine learning (e.g. Evgeniou et al., 2000). This GLM is trained on the training set, and compared against the test data in Fig. 4 and quantified in Table 2. The GLM predictions are generally good, but as before, the performance deteriorates in data sparse regions, for example near $\lambda/L = 1.5 \sim 2.0$ for Fr=0.3.

For the physics-based IM we use the linear version of the Large Amplitude Motion Prediction (LAMP-I) program (Lin and Yue, 1990), a time-domain three-dimensional potential flow panel method. The choice of LAMP-I as IM is based on the speed and stability of the linear potential-flow code. As in the previous application, all the model/computational parameters in LAMP-I are fixed, and the resulting predictions are shown in Fig. 4. As seen in the figure and in Table 2, the LAMP-I heave predictions are fairly accurate with RMSE=5.4% of the maximum response amplitude. However, the pitch accuracy is lower, particularly for the high speed case.

For the PBLM, we construct the phase shifted basis $\boldsymbol{\rho}$ of Eq. 13 using LAMP-I as the IM. The basis is incorporated into the GLM through the weighted model of Eq. 15 and retrained on the training data. The PBLM results are shown in Fig. 4 and the RMSE and eDOF summarized in Table 2. PBLM using a simple linear IM and 20 training points produce predictions with error levels below 3%, matching the performance of the time-domain RANS predictions of Weymouth et al. (2007). Moreover, the PBLM achieves this using only 10 degrees of freedom as opposed to the 16 degrees of freedom used by the GLM, whose predictions have around twice the error. Note that the inaccuracy of the IM in the high speed pitch case is not propagated to the PBLM response. As discussed in the introduction, an inaccurate IM is simply ignored by the learning model and the

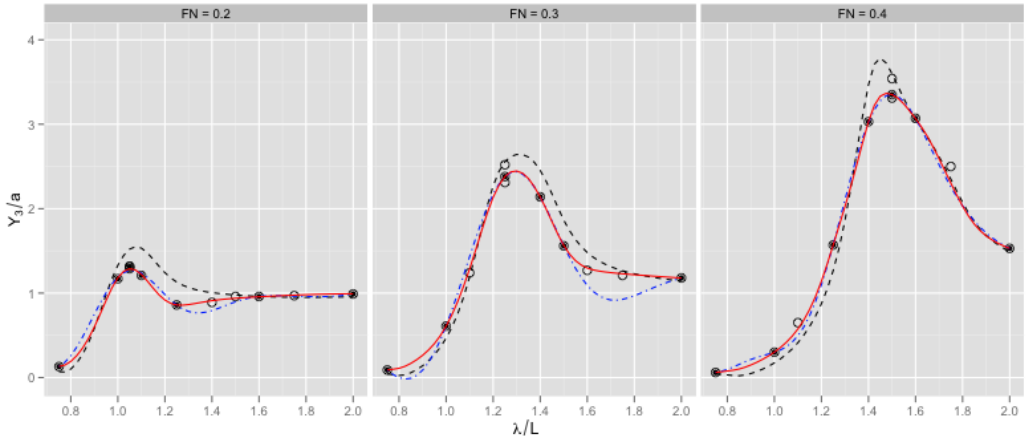

(a) Heave

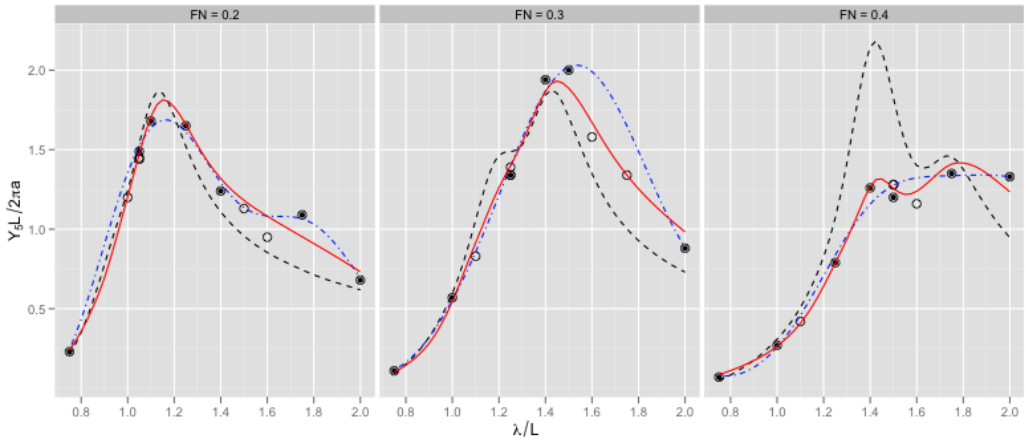

(b) Pitch

Figure 4: Heave and pitch response amplitude operators for the Wigley hull over a range of incident wavelengths and ship Froude numbers. Symbols denote the experimental measurements: • are the training data, ∘ are the test data. Line denote the prediction methods being compared: – · – is the GLM (using Gaussian regularization network); – – is the IM (using LAMP-I); — is the PBLM incorporating the IM in the GLM. The PBLM shows significantly increased accuracy away from training data, and correctly ignores the inaccurate LAMP-I pitch predictions for $Fr = 0.4$.

	Training points	GLM	PBLM
Ideal Training Set	20	0.042	0.025
No peak data	17	0.128	0.028
No $Fr = 0.3$ data	14	0.131	0.027

Table 3: Dependence of the Test RMSE on the training data set for the Wigley hull heave RAO predictions. The PBLM demonstrates a robust error level, $< 3\%$ despite the removal of critical training data.

resulted PBLM error matches that of the GLM in this region. Of more concern is the large variance in the IM for this test case which has caused some waviness in the PBLM. Low-pass filtering the IM prediction alleviates the waviness in the PBLM prediction but does not significantly change the error levels.

To explicitly demonstrate the data-dependence of the GLM predictions and the improvement in the PBLM approach, we repeat the heave predictions with two modified training sets. We note that our training set for the heave case is nearly ideal, spanning all three Froude numbers and includes all of the response peaks. While this yields accurate predictions, the situation is somewhat artificial. To evaluate the robustness of the different predictions with more limited training data, we show in Fig. 5 and Tab. 3 results for the same problem but now withholding from the training set: (a) peak response values; and (b) data for the intermediate Froude number. In case (a), seemingly crucial data is omitted; while in case (b), the total number of training points is reduced to only 14. In both cases, the PBLM performance is nearly unchanged with RMSE below 3%. In contrast, the GLM performance is substantially degraded with more than three times the error. It is possible that a more specialized GLM could be more appropriate for these extremely sparse data sets. However, the introduction of the IM and the physics-based basis $\boldsymbol{\rho}$ has eliminated the need for a more advanced learning method in these data-sparse cases.

3.3. Predictions of Breaking Bow Waves

Finally, we consider the prediction of the nonlinear bow waves of a theoretical hull form based on the 5415 with a sharp bow. This example is motivated by the desire to apply fast 2D+T modeling to the prediction of breaking bow waves and to demonstrate the ability of PBLM methods on a design-stage problem with no available experimental data.

The data for training and testing the learning models is obtained in this case from full-3D high-fidelity CFD simulations using the cartesian-grid Boundary Data Immersion Method (Weymouth and Yue, 2011) coupled with a conservative Volume of Fluid method (Weymouth and Yue, 2010). Simulations are run over three B/L slenderness values and three $F_{2D} = Fr\sqrt{\frac{B}{L}}$ speeds. This scaling ensures that the speeds span the range of different types of bow wave generation for each slenderness value. The low speed gives of a non-breaking bow wave, the intermediate speed gives a gentle spilling breaker, and the high speed produces a plunging breaker. As an example, the ship hull and the 3D CFD simulation prediction for the high speed, high thickness result is shown in Fig. 6. The simulations are processed to extract the height of the wave crest as a function of relative

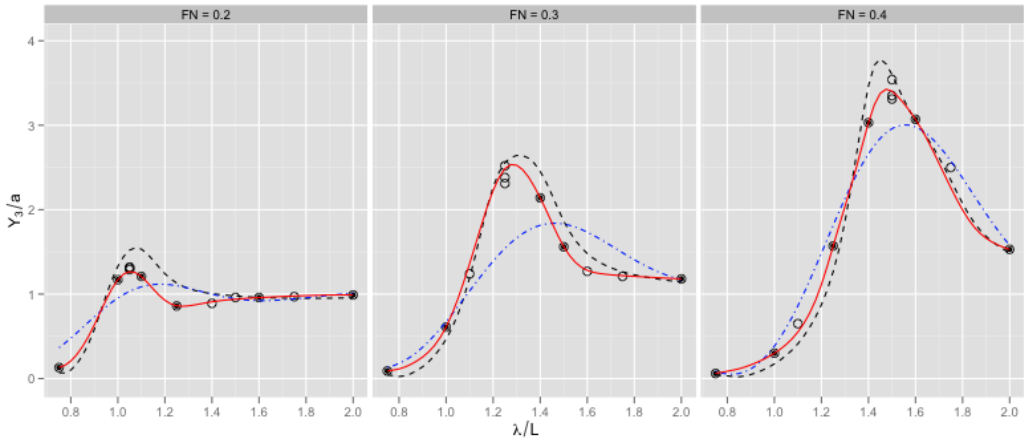

(a) No peak training data

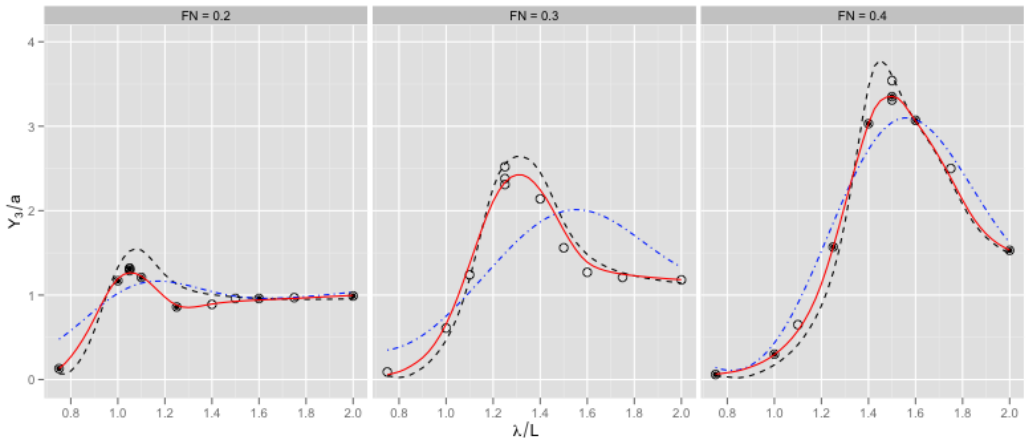

(b) No $Fr = 0.3$ training data

Figure 5: Heave RAO predictions for the Wigley hull using a reduced training set. The symbols and lines designations are as in Fig. 4. Relative to Fig. 4, the training set is reduced by (a) removing the peak response points; and (b) omitting all $Fr = 0.3$ data. In both cases, the PBLM prediction maintains its accuracy despite the removal of data; while the GLM performance is significantly affected.

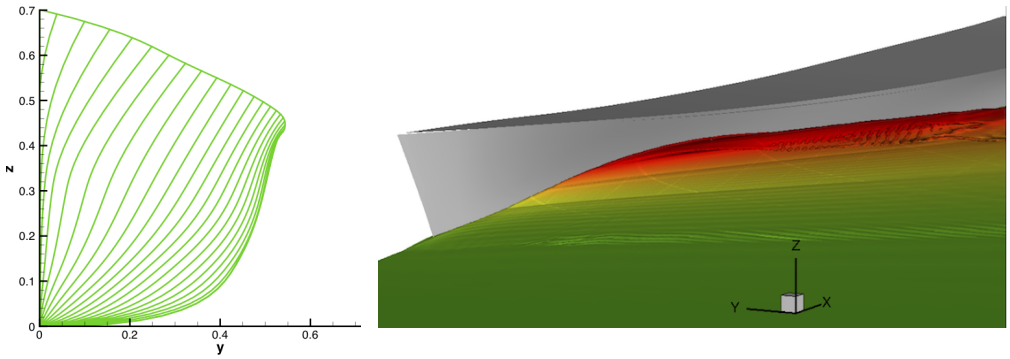

(a) Test geometry station lines

(b) 3D simulation $F_{2D} = 0.262, B/L = 0.12$

Figure 6: Geometry and sample 3D simulation for the bow wave test case. The geometry section lines (a) are based on a 5415 hull modified to have a very fine entrance angle. The full 3D CFD simulation (b) shows the hull in grey and the free surface colored by elevation. For this case ($F_{2D} = 0.262, B/L = 0.12$) there is a large plunging breaker at the bow.

position down the length of the bow (x/L). The data is shown as the symbols in Fig. 8 where it can be seen that the extraction function introduces a level of background noise.

For the GLM in this case, we use the Support Vector Machine Regression (Vapnik, 1995). The training data supplied to the GLM include the high and low speed data, with the intermediate speed data withheld for testing. The GLM predictions are shown in Fig. 8 and summarized by B/L-value in Tab. 4. As with the other GLMs, Support Vector Machine Regression shows good ability to fit the training data and correctly ignores the noise. However its generalization to the unseen Froude number test data is fairly poor, especially for the largest B/L. In order for a generic learning model to make accurate predictions for this case, many more 3D simulations would need to be run across the range of F_{2D}, at a significant increase of overall computational cost.

For the physics-based IM, we use a Cartesian-grid 2D+T model. The 2D+T approach is formally based on slender body theory assuming that changes in the longitudinal direction are small compared to changes in the transverse directions (e.g. Fontaine and Cointe, 1997; Fontaine and Tulin, 2001; Weymouth et al., 2006). In this approach, the 2D+T representation of the ship geometry is that of a 2D flexible wavemaker, whose instantaneous profile matches the section lines of the 3D hull. The successive 2D waves generated by this wave-maker correspond to the divergent waves of the 3D ship in the limit of high ship speed and slenderness. As there is no longitudinal length or velocity in the wavemaker flow, the 2D+T 'speed' is set through the time $T = L/U$ it takes for the wavemaker to trace the sectional shape of the vessel. Thus the appropriate non-dimensional scaling is $F_{2D} = \sqrt{\frac{B}{gT^2}}$. Although theoretically only valid for a slender ship moving at high speed, approximating the actual three-dimensional flow by a two-dimensional, time evolving flow reduces the computational cost by orders of magnitude, making it attractive for use as the IM in a PBLM.

The 2D+T $f_{\mathcal{I}}$ predictions are obtained using the same Cartesian-grid method and over the same F_{2D} values as in the 3D tests. While essentially the same code, the 2D+T predictions are complete in around a minute each on a desktop, as opposed to the 3D

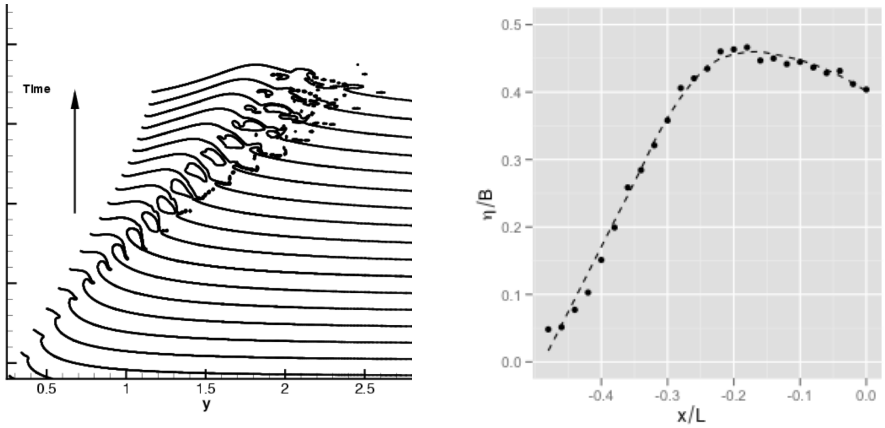

(a) 2D+T waterfall plot

(b) Wave crest height

Figure 7: 2D+T results for the high speed case $F_{2D} = 0.262$. (a) Waterfall plot generated by overlaying the predicted free-surface elevations at different times, shifted vertically for ease of viewing. (b) The extracted wave crest heights, with symbols for the raw data and the dashed line for the smoothed IM prediction used in $\boldsymbol{\rho}$.

	2D+T	GLM	PBLM
$B/L = 0.03$	0.043	0.106	0.031
$B/L = 0.06$	0.081	0.098	0.055
$B/L = 0.12$	0.191	0.302	0.112

Table 4: RMSE, scaled by the maximum response of $\eta/(BF_{2D}) = 2.11$, on the test set for the breaking wave crest elevation predictions (for the high speed $F_{2D} = 0.191$ case). Despite higher error in the 2D+T IM at greater B/L, PBLM outperforms GLM on the test data.

runs which take $O(10^4)$ CPU hours at a high performance computing facility. The result for the high speed case is shown in Fig. 7(a). The 2D+T model predicts a plunging breaking wave for the high speed test, the same as the 3D CFD prediction. To quantify the comparison, we extract the location of the wave crest as shown in Fig. 7(b). This the extraction introduces noise into the 2D+T results. As noted in the introduction, it is important to remove this noise with a low-pass filter to avoid the noise propagating up to the PBLM predictions. The final IM prediction is shown (as dashed lines) in Fig. 7(b) and Fig. 8. Note that the 2D+T prediction has no dependence on B/L (the curves are the same for each column of Fig. 8) and are only accurate for the slender high speed ship, as summarized in Tab. 4.

For the PBLM, we construct $\boldsymbol{\rho}$ using Eq. 13 with the 2D+T IM, incorporating the basis into the GLM using the weighted model Eq. 15, and retrained on the training data. The resulting PBLM predictions are plotted in Fig. 8 and summarized for different B/L in Tab. 4. The results show that the trained PBLM makes excellent quantitative predictions for $B/L = 0.03$ where the 2D+T model gives a good estimate of the waves at all speeds. As the ship thickness increases so does the error in the 2D+T predictions, but the 2D+T IM still provides information on the dependence of the bow waves on Froude number. Complementing this information with the high and low Froude number data at

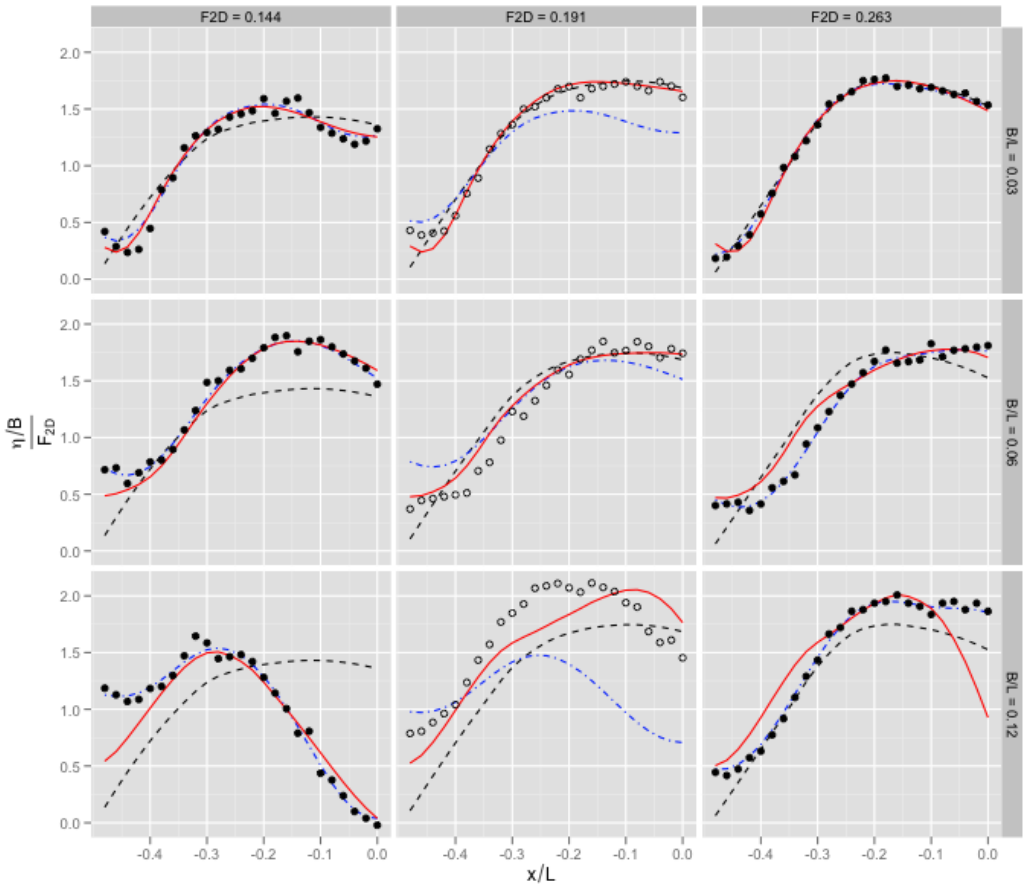

Figure 8: Data and predictions for the wave crest elevation of breaking bow waves at a function of 2D Froude number (F_{2D}) and ship thickness (B/L). Symbols denote measured values from the 3D Cartesian-grid CFD: • are the training data, ∘ are the test data. Lines denote the prediction methods: −− is the 2D+T IM prediction; – · – the GLM prediction; and — the PBLM prediction. PBLM enables predictions across the Froude number range with higher accuracy than the GLM even for larger B/L values.

each B/L enables the PBLM to reduce the GLM errors by more than 60% on the thickest hull. Thus, using PBLM with the fast 2D+T IM, we are able to generalize relatively few 3D data points across the Froude number range to generate reliable predictions that may be used, in this case, to evaluate the effect of ship geometry and speed on nonlinear breaking bow waves, at low computational cost.

4. Discussion and Conclusions

We present a general method for constructing physics-based learning models (PBLM) which incorporate a physics-based basis to increase predictive accuracy and decrease training data dependence. The method requires the use of a generic learning model (GLM), a fast physics-based intermediate model (IM), and a small set of high-quality experimental or computational training data from the physical system. We demonstrate the effectiveness of PBLM for three different ship hydrodynamics problems. To highlight the generality and versatility of the new approach, we used different sources for the training data (from experiments and computations), different generic learning models, and different intermediate models in these applications. In each case, PBLM obtains greatly improved prediction accuracy and robustness over GLM, and is orders of magnitude faster than high-fidelity CFD.

In these examples we have presented only the most basic PBLM approach, using only one source of data and a single IM in each case. PBLM is directly applicable to combinations of different data (say experiments plus CFD) and multiple intermediate models (numerical, analytic or empirical). For example, the addition of an IM which described the variation in B/L for the bow wave predictions of §3.3 would enable even higher accuracy predictions across the design space. Additionally, we have only presented one method of generating the physical basis $\boldsymbol{\rho}$. Techniques to establish a complete orthogonal basis, such as PCA, and use this in a PBLM framework are currently in development.

Increased performance demands on naval and marine platforms have led to modern designs that require very large numbers of design evaluations. This requirement is met by tank and field experiments and high-fidelity simulations, which are expensive; complemented by approximate design and simulation tools, which do not capture all aspects of the problems or to sufficient accuracy. The proposed PBLM provides a general and powerful framework which combines the strength of these separate approaches to obtain significantly more useful predictions than either in isolation. This efficacy of PBLM increases the value of both high-fidelity expensive data and approximate fast tools and should prove useful in all stages of modern analysis and design.

Acknowledgements We would like to acknowledge Sheguang Zhang, Claudio Cairoli, and Thomas Fu for help supplying the intermediate and experimental sources. This research was supported financially by grants from the Office of Naval Research and the computational resources were provided through a Challenge Project grant from the DoD High Performance Computing Modernization Office.

References

N.M. Alexandrov, JE Dennis, R.M. Lewis, and V. Torczon. A trust-region framework for managing the use of approximation models in optimization. *Structural and Multidisciplinary Optimization*, 15(1): 16–23, 1998.

N.M. Alexandrov, R.M. Lewis, C.R. Gumbert, L.L. Green, and P.A. Newman. Approximation and model management in aerodynamic optimization with variable-fidelity models. *Journal of Aircraft*, 38(6):1093–1101, 2001.

Christopher J.C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2:121–167, 1998.

Pablo M. Carrica, Robert V. Wilson, Ralph W. Noack, and Fred Stern. Ship motions using single-phase level set with dynamic overset grids. *COMPUTERS & FLUIDS*, 36(9):1415–1433, NOV 2007. ISSN 0045-7930. doi: 10.1016/j.compfluid.2007.01.007.

G. Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367, 2003.

Theodoros Evgeniou, Massimiliano Pontil, and Tomaso Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13, 2000.

Wil Faller, David Hess, and Thomas Fu. Simulation based design: A real-time approach using recursive nueral networks. *AIAA*, 0911:1–13, 2005.

Wil Faller, David Hess, and Thomas Fu. Real-time simulation based design part ii: Changes in hull geometry. *AIAA*, 1485:1–16, 2006.

E. Fontaine and R. Cointe. A slender body aproach to nonlinear bow waves. *Phil. Trans. Royal Society of London*, A 355:565–574, 1997.

E. Fontaine and M. P. Tulin. On the prediction of nonlinear free-surface flows past slender hulls using 2d+t theory. *Ship Tech. Res.*, 48, 2001.

W. Froude. *The Papers of William Froude*. London: Institution of Naval Architects, 1955.

F. Girosi, M. Jones, and T. Poggio. Regularization theory and neural networks architectures. *Neural computation*, 7(2):219–269, 1995.

G.H. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, pages 215–223, 1979.

P.S. Granville. Drag and turbulent boundary layer of flat plates at low reynolds numbers. *Journal of Ship Research*, 21, 1977.

David E. Hess, Wil E. Faller, Thomas C. Fu, and Edward S. Ammeen. Improved simulation of ship maneuvers using recursive neural networks. *AIAA*, 1481, 2006.

JL Hess and AMO Smith. Calculation of nonlifting potential flow about arbitrary three-dimensional bodies. *JOURNAL OF SHIP RESEARCH*, 8(2):22–44, 1964.

J. M. J. Journee. Experiments and calculations on four wigley hullforms. Technical Report 909, Delft University of Technology, Ship Hydrodynamic Laboratory,, February 1992.

Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(Series D):35–45, 1960.

W.M. Lin and D.K.P Yue. Numerical solutions for large-amplitude ship motions in the time-domain. In *Proceedings of the 18th Symposium on Naval Hydrodynamics*, pages 41–66, The University of Michigan, U.S.A., 1990.

B. Scholkopf, A.J. Smola, R. Williamson, and P. Bartlett. New support vector algorithms. *Neural computation*, 12(5):1207–1245, 2000.

Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization and Beyond*. MIT Press, 2001.

J.R. Shearer and J.J. Cross. The experimental determination of the components of ship resistance for a mathematical model. 1965.

Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1995.

G. Weymouth, K. Hendrickson, D.K.P. Yue, T. O'Shea, D. Dommermuth, P. Adams, and M. Valenciano. Modeling breaking ship waves for design and analysis of naval vessels. *hpcmp-ugc*, pages 440–445, 2007.

Gabriel D. Weymouth and Dick K.-P. Yue. Boundary data immersion method for cartesian-grid simulations of fluid-body interaction problems. *Journal of Computational Physics*, 2011.

Gabriel D. Weymouth, Douglas G. Dommermuth, Kelli Hendrickson, and Dick K.-P. Yue. Advancements in cartesian-grid methods for computational ship hydrodynamics. In *26th Symposium on Naval Hydrodynamics*, 2006.

G.D. Weymouth and Dick K.-P. Yue. Conservative volume-of-fluid method for free-surface simulations

on cartesian-grids. *Journal of Computational Physics*, 229(8):2853 – 2865, 2010. ISSN 0021-9991. doi: DOI: 10.1016/j.jcp.2009.12.018.

GD Weymouth, RV Wilson, and F Stern. Rans computational fluid dynamics predictions of pitch and heave ship motions in head seas. *JOURNAL OF SHIP RESEARCH*, 49(2):80–97, JUN 2005. ISSN 0022-4502.

S. N. Wood. Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65:95–114, 2003.

S. N. Wood. Fast stable direct fitting and smoothness selection for generalized additive models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70:495–518, 2008.